What is claimed is:

1. A method of reading data comprising the steps of:

receiving a request for a data block stored within a stripe of erasure coded data, the stripe of erasure coded data stored across a plurality of storage devices which include a target storage device that holds the data block;

sending read messages to the storage devices; and

receiving a reply message from each of at least a quorum of the storage devices, the reply message from the target storage device including the data block, the quorum meeting a quorum condition of a number such that any two selections of the number of stripe blocks intersect in a minimum number of the stripe blocks needed to decode the stripe.

2. The method of claim 1 wherein the stripe blocks comprise a first number of data blocks and a second number of parity blocks.

3. The method of claim 1 wherein the read message that is sent to the target storage device includes an indicator.

4. The method of claim 1 wherein each of the reply messages of the quorum indicate that there is no pending write for the stripe block stored on the storage device.

5. The method of claim 1 wherein the reply messages of the quorum include validation timestamps that match.

6. The method of claim 1 wherein the stripe of erasure coded data was previously stored using a technique of striping.

7. The method of claim 1 wherein the stripe of erasure coded data was not previously stored using a technique of striping.

8. A method of reading data comprising the steps of:

receiving a request for a data block stored within a stripe of erasure coded data, the stripe of erasure coded data stored across a plurality of storage devices which include a target storage device that holds the data block, the stripe comprising stripe blocks;

sending read messages to the storage devices, the read message sent to the target storage device including an indicator; and

receiving a reply message from each of at least a quorum of the storage devices which indicate that there is no pending write for the stripe block stored on the storage device, the reply messages including validation timestamps which match, the reply message from the target storage device including the data block, the quorum meeting a quorum condition of a number such that any two selections of the number of stripe blocks intersect in the minimum number of the stripe blocks needed to decode the stripe.

9. A method of writing data comprising the steps of:

receiving a data block for storage within a stripe of erasure coded data, the stripe comprising stripe blocks;

sending a query message to each of a plurality of storage devices upon which the stripe of erasure coded data is stored;

receiving a query reply message from each of at least a first quorum of the storage devices;

sending a modify message to each of the storage devices; and

receiving a write reply message from each of at least a second quorum of the storage devices, the first and second quorums each meeting a quorum condition of a number such that any two selections of the number of the stripe blocks intersect in a minimum number of the stripe blocks needed to decode the stripe.

10. The method of claim 9 wherein the stripe blocks comprise a first number of data blocks and a second number of parity blocks.

11. The method of claim 9 wherein the query message sent to the storage device upon which the data block is to be stored includes an indicator.

12. The method of claim 9 wherein the query messages include a timestamp indicating a current time.

13. The method of claim 12 wherein the reply messages of the first quorum each indicate that the timestamp is later than a pending write timestamp.

14. The method of claim 12 wherein the reply messages of the first quorum each indicate that the timestamp is later than a validation timestamp for a previous version of the data block.

15. The method of claim 14 wherein the query reply message from the storage device upon which the data block is to be stored includes the validation timestamp and the previous version of the data block.

16. The method of claim 15 wherein each of the modify messages include the timestamp and the validation timestamp.

17. The method of claim 16 wherein the modify message sent to the storage device upon which the data block is to be stored includes the data block.

18. The method of claim 17 wherein the modify messages sent to the storage devices which hold parity blocks include the previous version of the data block and the data block.

19. The method of claim 18 wherein the write reply messages from the second quorum indicate that the validation timestamp equals a maximum timestamp for the stripe block stored on the storage device.

20. The method of claim 18 wherein the write reply messages from the second quorum indicate that the timestamp is no earlier than the pending write timestamp.

21. The method of claim 18 wherein the write reply message from the storage device which stores the data block indicates that the data block was stored

successfully.

22. The method of claim 17 wherein the modify messages sent to the storage devices which hold parity blocks include a coded block which represents the previous version of the data block and the data block.

23. The method of claim 9 wherein the stripe of erasure coded data was previously stored using a technique of striping.

24. The method of claim 9 wherein the stripe of erasure coded data was not previously stored using a technique of striping.

25. A method of writing data comprising the steps of:

receiving a particular data block for storage within a stripe of erasure coded data, the stripe comprising stripe blocks which comprise a first number of data blocks and a second number of parity blocks;

sending a query message including a timestamp indicating a current time to each of a plurality of storage devices upon which the stripe of erasure coded data is stored, the query message sent to a target storage device upon which the particular data block is to be stored including an indicator;

receiving a query reply message from each of at least a first quorum of the storage devices indicating that the timestamp is later than a pending write timestamp and that the timestamp is later than a validation timestamp for the earlier version of the stripe block, the query reply message from the target storage device including the validation timestamp and a previous version of the particular data block;

sending a modify message to each of the storage devices including the timestamp and the validation timestamp, the modify message sent to the storage device upon which the block of data is to be stored including the particular data block, the modify messages sent to the storage devices which hold the parity blocks including information used to update the parity blocks; and

receiving a write reply message from each of at least a second quorum

of the storage devices indicating that the validation timestamp equals a maximum timestamp for the stripe block stored on the storage device, and that the timestamp is no earlier than the pending write timestamp, the first and second quorums each meeting a quorum condition of a number such that any two selections of the number of the stripe blocks intersect in a minimum number of the stripe blocks needed to decode the stripe.

26. The method of claim 25 wherein the information used to update the parity blocks comprises a coded block which represents the previous version of the particular data block and the particular data block.

27. The method of claim 25 wherein the information used to update the parity blocks comprises the previous version of the particular data block and the particular data block.

28. A computer readable memory comprising computer code for implementing a method of reading data, the method of reading the data comprising the steps of:

receiving a request for a data block stored within a stripe of erasure coded data, the stripe of erasure coded data stored across a plurality of storage devices which include a target storage device that holds the data block;

sending read messages to the storage devices; and

receiving at least a quorum of reply messages from the storage devices, the reply message from the target storage device including the data block, the quorum meeting a quorum condition of a number such that any two selections of the number of stripe blocks intersect in a minimum number of the stripe blocks needed to decode the stripe.

29. The computer readable memory of claim 28 wherein the stripe blocks comprise a first number of data blocks and a second number of parity blocks.

30. The computer readable memory of claim 28 wherein the read message that is sent to the target storage device includes an indicator.

31. The computer readable memory of claim 28 wherein the reply messages of the quorum indicate that there is no pending write for the stripe block stored on the storage device.

32. The computer readable memory of claim 28 wherein the reply messages of the quorum include validation timestamps that match.

33. The computer readable memory of claim 28 wherein the stripe of erasure coded data was previously stored using a technique of striping.

34. The computer readable memory of claim 28 wherein the stripe of erasure coded data was not previously stored using a technique of striping.

35. A computer readable memory comprising computer code for implementing a method of writing data, the method of writing the data comprising the steps of:

receiving a data block for storage within a stripe of erasure coded data, the stripe comprising stripe blocks;

sending a query message to each of a plurality of storage devices upon which the stripe of erasure coded data is stored;

receiving a query reply message from each of at least a first quorum of the storage devices;

sending a modify message to each of the storage; and

receiving a write reply message from each of at least a second quorum of the storage devices, the first and second quorums each meeting a quorum condition of a number such that any two selections of the number of the stripe blocks intersect in a minimum number of the stripe blocks needed to decode the stripe.

36. The computer readable memory of claim 35 wherein the stripe blocks comprise a first number of data blocks and a second number of parity blocks.

37. The computer readable memory of claim 35 wherein the query message sent to the storage device upon which the data block is to be stored includes an indicator.

38. The computer readable memory of claim 35 wherein the query messages include a timestamp indicating a current time.

39. The computer readable memory of claim 38 wherein the reply messages of the first quorum each indicate that the timestamp is later than a pending write timestamp.

40. The computer readable memory of claim 38 wherein the reply messages of the first quorum each indicate that the timestamp is later than a validation timestamp for a previous version of the data block.

41. The computer readable memory of claim 40 wherein the query reply message from the storage device upon which the data block is to be stored includes the validation timestamp and the previous version of the data block.

42. The computer readable memory of claim 41 wherein each of the modify messages include the timestamp and the validation timestamp.

43. The computer readable memory of claim 42 wherein the modify message sent to the storage device upon which the data block is to be stored includes the data block.

44. The computer readable memory of claim 43 wherein the modify messages sent to the storage devices which hold parity blocks include the previous version of the data block and the data block.

45. The computer readable memory of claim 44 wherein the write reply messages from the second quorum indicate that the validation timestamp equals a maximum timestamp for the stripe block stored on the storage device.

46. The computer readable memory of claim 44 wherein the write reply messages from the second quorum indicate that the timestamp is no earlier than the pending write timestamp.

47. The computer readable memory of claim 44 wherein the write reply message from the storage device which stores the data block indicates that the data block was stored successfully.

48. The computer readable memory of claim 43 wherein the modify messages sent to the storage devices which hold parity blocks include a coded block which represents the previous version of the data block and the data block.

49. The computer readable memory of claim 35 wherein the stripe of erasure coded data was previously stored using a technique of striping.

50. The computer readable memory of claim 35 wherein the stripe of erasure coded data was not previously stored using a technique of striping.